Hardness of Robust Network Design*

C. Chekuri and F. B. Shepherd
Lucent Bell Labs, 600 Mountain Ave, Murray Hill, New Jersey 07974

G. Oriolo
Dipartimento di Ingegneria dell'Impresa, via del Politecnico 1, 00133, Rome, Italy

M. G. Scutellá
Dipartimento di Informatica, Largo Bruno Pontecorvo 3, 56127, Pisa, Italy

The authors settle the complexity status of the robust network design problem in undirected graphs. The fact that the flow-cut gap in general graphs can be large, poses some difficulty in establishing a hardness result. Instead, the authors introduce a single-source version of the problem where the flow-cut gap is known to be one. They then show that this restricted problem is coNP-Hard. This version also captures, as special cases, the fractional relaxations of several problems including the spanning tree problem, the Steiner tree problem, and the shortest path problem. © 2007 Wiley Periodicals, Inc. NETWORKS, Vol. 50(1), 50–54 2007

Keywords: network design; single-source hose model; robust optimization

1. INTRODUCTION

A crucial assumption in many network design problems is that of knowing the traffic demands in advance. Unfortunately, measuring and predicting traffic demands are difficult problems. Moreover, in several applications, communication patterns change over time, and therefore we are not given a *single* static traffic demand, but instead a set of *nonsimultaneous* traffic demands. The network should be able to support any traffic demand that is from the given set.

Received June 2005; accepted June 2006
Correspondence to: M.G. Scutellá; e-mail: scut@di.unipi.it
*A preliminary version has appeared in Proc INOC, (2005), 455–461.
Contract grant sponsor: ONR basic research; Contract grant number: N00014-04-M-0042
Contract grant sponsor: Italian Research Project Cofinanziato 2003 "Ottimizzazione, simulazione e complessitá nel progetto e gestione di reti di telecomunicazione"
DOI 10.1002/net.20165
Published online in Wiley InterScience (www.interscience.wiley.com).

In this paper, we consider the following *robust network design problem*. We are given a graph $G = (V, E)$ on n nodes and a convex body P_D of $n \times n$ non-negative "demand" matrices. We work primarily with undirected graphs, although our main result holds also in the directed setting. We denote by $\mathcal{U}(P_D)$ the set of all *capacity reservations* in G that *support* every demand matrix $D \in P_D$. We make this more precise now.

An instance will be called *oriented* if a demand between a pair of nodes i, j is distinguished by whether it *starts* at node i, or at node j. Otherwise, it is called *unoriented*; in this case, D is a symmetric matrix where, for each i, j, the entries $D_{ij} = D_{ji}$ represent the same single demand between i, j. For instances where G is a directed graph, all demands are necessarily oriented since they must be routed on directed paths from i to j. However, in undirected graphs, demands may be oriented or unoriented. Thus, we allow demand matrices D to be asymmetric even in undirected graphs, and, in this case, D_{ij} and D_{ji} identify distinct demands starting from i and j, respectively. We work now in this oriented setting as it is more general; for instance, one can work with lower triangular matrices to mimic unoriented demands.

We now define the region $\mathcal{U}(P_D)$. Let $\mathcal{P}_{ij}$ denote the finite set of simple paths between i and j. (In principle, we could also define P_{ij} to only be a subset of such paths which are "feasible" for the demand from i to j.) An edge capacity reservation $u : E \to \mathbf{R}_+$ is in $\mathcal{U}(P_D)$ (for oriented demands) if for each $D \in P_D$ and for each ordered pair ij, there is an assignment f_P^{ij} of flows to paths P in G such that for each ij: $\sum_{P \in \mathcal{P}_{ij}} f_P^{ij} = D_{ij}$ and for each edge $e \in E$: $\sum_{ij} \sum_{P \in \mathcal{P}_{ij}: e \in P} f_P^{ij} \le u_e$. We call the assignment f a *feasible flow* for the demands D. Note that $\mathcal{U}(P_D)$ is convex. Given a cost vector c, our aim is to find a vector $u \in \mathcal{U}(P_D)$ that minimizes cu. In this note, we settle the complexity of robust network design, by establishing its coNP-hardness, even when P_D is given by a set of explicit constraints whose size is polynomial in n. To our knowledge, this problem was first suggested in [4]; moreover, an $O(\log n)$ approximation can be obtained (Gupta, personal communication, 2004).

We emphasize that, in our model, the routing can vary as the traffic demand changes. The same hypothesis has been implicitly considered for a robust network design problem in a recent paper [22], where the objective is to minimize an appropriate measure of the set of the unsatisfied demands; the authors propose a cutting-plane algorithm for the special case in which P_D is the convex hull of a discrete number of scenarios.

Related work has addressed *oblivious* routing, which is sometimes also referred to as *stable* routing. The routing is oblivious in the sense that a common *routing template* is used for each relevant traffic demand matrix. More precisely, for each ordered pair of nodes ij (or unordered pair $\{i,j\}$ in some contexts), the routing specifies a unit flow: if D_{ij} demand needs to be routed for some traffic matrix $D \in P_D$, then it simply scales the unit flow by D_{ij}. The *robust network design problem with oblivious routing* is considered in [4], where a column-generation algorithm is studied. An introduction to the use of polytopes for modeling traffic demand, under an oblivious routing template, is given in [5]. The authors also suggest a period-dependent routing as an intermediate strategy between oblivious routing and fully dynamic routing, owing to the inherent impracticalities to implement the latter.

Another line of work seeks oblivious routings that minimize congestion. More precisely, suppose we are given network G with a capacity reservation $u : E \to \mathbf{R}_+$. A (not necessarily feasible) flow is said to have *congestion* α if it becomes feasible after scaling up each capacity by a factor of α. A *congestion bound* for a routing template (i.e., for an oblivious routing) is a value β such that if a demand matrix D can be routed by some flow with congestion α, then routing D using the template has a congestion of at most $\alpha\beta$. For undirected graphs, Räcke [24] showed the existence of oblivious routings with a congestion bound that is polylogarithmic in n. In [2, 3] algorithms are given (both for directed and undirected graphs) to find oblivious routings with optimal congestion bounds. One may formulate this problem of finding an oblivious routing, with a minimum congestion bound, as a robust network optimization problem as follows. Let P_D be the set of traffic matrices that can be routed in G with congestion at most 1. We may restrict to these matrices since for any value $a > 1$, if a flow f routes a matrix D with congestion β, then the flow $(1/a)f$ routes the matrix $(1/a)D$ with congestion β. Hence a routing template has a congestion bound of β if routing every demand matrix $D \in P_D$, using the template, has a congestion of at most β. Thus the problem is to find a minimum β for which there is a routing template such that each demand matrix in P_D is routed obliviously using capacity at most βu_e for each edge e. In [3], a polynomial time algorithm based on the ellipsoid algorithm is given for this problem. In [2] another algorithm is given by formulating the problem as a compact linear program. The latter paper also considers the polytope P_D consisting of all matrices which are "near" to a given target matrix. We note that the ellipsoid algorithm in [3] also generalizes to the setting where one seeks a oblivious routing with minimum congestion for demand matrices belonging to a given polytope P_D. This is essentially the separation problem for the robust network design with oblivious routing.

One well-studied special case of robust network design with oblivious routing is where traffic is assumed to obey the so-called *hose model*, introduced by Fingerhut et al. [10] and Duffield et al. [7]. In the *symmetric marginals model*, each node i specifies a *marginal* traffic amount b_i (usually an integer), and a *valid* traffic matrix D is an assignment of unoriented demands D_{ij} that respects the cumulative upper bounds $\{b_i\}$: for each i, we have $\sum_j D_{ij} \le b_i$, where $D_{ij} = D_{ji}$ represent the demand between i and j.

In the *asymmetric marginals* version of the hose model, each node i has two upper bounds, b_i^+, b_i^-, which represent upper bounds on the amount of flow that the node i can send and receive, respectively. In other words, for a given bound vector b, the region P_D, denoted by $P(G,b)$, is defined as:

$$P(G,b) := \Big\{ D \ge 0 : \sum_j D_{ij} \le b_i^+ \text{ and } \sum_j D_{ji} \le b_i^- \text{ for each node } i \Big\}. \quad (1)$$

In the asymmetric marginal case, the network design problem is known to be coNP-hard in directed graphs [14]. The complexity of the problem in undirected graphs has not been established, even in the symmetric marginals case when all b_is are 1.

A related problem that has received particular attention is the so-called *(asymmetric) virtual private network design problem* (VPN problem). In this problem, the region P_D is defined as in the asymmetric marginal case of the hose model. The routing is required to be oblivious, and furthermore it is restricted to use a single path for each pair of nodes. The asymmetric VPN problem is coNP-hard, even when we restrict to tree reservations [14]. Nevertheless, there are cases where finding an optimal tree reservation is easy, namely: (1) the asymmetric *balanced* case, i.e., when $\Sigma_{v\in T} b_i^+ = \Sigma_{v\in T} b_i^-$ [15]; (2) the symmetric case [14]. Constant factor approximation algorithms for both the asymmetric and the symmetric VPN problems can be found in [8, 13–15, 19], see [8] for a survey. In [1] a compact linear mixed-integer programming formulation for the asymmetric problem is given; it is also shown that this compact formulation allows to solve medium-to-large-size instances to optimality with commercial solvers, while a branch-and-price and cutting plane algorithm allows to tackle larger instances.

In [9] a cutting plane algorithm is devised to obtain empirical evidence that allowing fractional routing templates leads to significant cost reductions for the hose model. In [23] several robust design problems are considered; a "robustness premium" is measured and empirical evidence based on equipment costs is rendered. In [20], an argument similar to one from [2] is used to obtain a compact LP formulation for maximizing throughput of oblivious routings for the class of hose demand matrices.

Confusion often occurs since instances may have oriented or unoriented demands, asymmetric or symmetric marginals, and flow vectors may be in directed or undirected graphs. It is worth keeping in mind that for directed graphs, or undirected graphs with asymmetric marginals, only oriented demands make sense. While our result is essentially about unoriented demands in undirected graphs, we found it simplest to present our proof in terms of oriented demands in a single-source flow problem in undirected graphs.

2. HARDNESS OF ROBUST NETWORK DESIGN

We show that the separation problem for $\mathcal{U}(P_D)$ is coNP-complete. From the equivalence of separation and optimization for linear programming [12], we obtain coNP-hardness of robust network design. The key to this proof is to consider a simpler class of robust design problems arising from single-source demand matrices. In a *single-source robust network design* instance, there is some *root node* $r \in V$ such that each matrix in P_D assigns positive demands only to pairs of the form ri, while $D_{ij} = 0$ if $r \neq i$.

$$\begin{aligned}\sum_{i \neq r} D_{ri} &\leq b_r \\ D_{ri} &\leq b_i \qquad \forall\, i \neq r \\ D_{ij} &= 0 \qquad r \neq i \\ D &\geq 0.\end{aligned}$$

We denote by $P(G, r, b)$ the corresponding polytope of demand matrices. We also denote by $\mathcal{U}(G, r, b)$ the region $\mathcal{U}(P(G, r, b))$. We recall that the *dominant* of a convex body P in $\mathbf{R}^n$ is the convex region $\{y \in \mathbf{R}^n : y \geq x \text{ for some } x \in P\}$.

We later need the following which is a simple exercise.

Fact 2.1. *For integer vector b, each extreme point of $P(G, r, b)$ is an integral demand matrix D.*

The single-source hose problem includes a number of interesting special cases. Suppose b is a 0,1 vector with $b_r = 1$ and $Z = \{i : b_i = 1\}$. Then, $\mathcal{U}(P_D)$ is the dominant of the fractional spanning tree polytope if $Z = V$, and is otherwise the dominant of the fractional Steiner tree polytope with terminal set Z. If $b_r = |V| - 1$, and each $b_i = 1$ if $i \neq r$, then $\mathcal{U}(P_D)$ is the dominant of the convex hull of path trees rooted at r (i.e., edge capacitated trees resulting from sending one unit of flow from each terminal to r). In fact, minimum cost reservations will then correspond to shortest path trees rooted at r. Note that the separation problem is easy in both the preceding cases, and indeed also when b is integer with b_r fixed. This is also true when G is directed, where the two extremes correspond to the minimum arborescence problem and shortest path routing trees.

A convex region P is a *blocking convex body* if it is equal to its dominant. Note that $\mathcal{U}(P_D)$ is a blocking polyhedron for every polyhedron P_D. Given $\alpha \geq 1$, an *α-separation algorithm* for a convex polyhedron P is an algorithm that does the following: given a vector z, it either declares that $z \notin P$ and exhibits a violated hyperplane (i.e., an inequality $ax \leq \gamma$ such that each element of P satisfies this inequality, but z does not), or it correctly declares that $\alpha z \in P$. We observe that the approximate separation algorithm might in some cases declare that $z \notin P$ even if $\alpha z \in P$. If $\alpha = 1$, we simply refer to this as a separation algorithm.

Let $G' = (V', E')$ be an undirected graph with a given nonnegative edge capacity function $u : E' \to \mathbf{R}^+$. A subset $S \subset V'$ is a *light set* if $|S| \leq |V'|/2$. The *expansion* of a set S is $exp(S) := \frac{u(\delta_{G'}(S))}{|S|}$, where $u(\delta_{G'}(S))$ is the sum of the capacities of the edges crossing the cut $(S, V' \setminus S)$. The *expansion* of G' is the minimum expansion of a light set: $exp(G') = \min_{S \subset V', |S| \leq |V'|/2} exp(S)$. It is well known [11, 21] that it is coNP-hard to determine the expansion of a graph. (Expansion is typically defined for uncapacitated graphs; [21] refers to the weighted expansion as *flux*.)

We will prove now that, for any integer $k \geq 1$, if there is a polynomial time $2\frac{k}{k+1}$-separation algorithm for the single-source hose problem, then there is a polynomial time algorithm which, given an input graph G', outputs YES if $exp(G') \geq 1$, NO if $exp(G') < 1/k$, and otherwise its output is not defined. Since it is coNP-complete to decide if the expansion of a given graph is at least 1 [11, 21] then, by establishing the above result, and taking $k = 1$, we obtain coNP-completeness of the separation problem for the single-source hose problem.

Let G' be a given edge weighted graph. From G' we obtain an instance of the single-source hose design problem as follows. Construct a graph $G = (V, E)$ by taking two copies G_1, G_2 of G'. For an edge $e \in G'$ with capacity u_e, we set the capacities of edges e_1 and e_2 equal to $u_e/2$, where e_1 and e_2 are the copies of e in G_1, G_2. For every node v in G' we add an edge of infinite capacity between the two copies v_1, v_2 of v. (A capacity of $|V'|$ would be sufficient but it is comforting to think of the capacity as infinite.) We add a root node r and add an edge of capacity 1 from r to each node in G_1. See Figure 1. Let x be the resulting edge capacity vector on G. We set $b_r = (|V| - 1)/2 = |V'|$ and $b_i = 1$ for each $i \neq r$.

Claim 2.2. *If $exp(G') \geq 1$, then $x \in \mathcal{U}(G, r, b)$. If there is some light subset S of G' such that $exp(S) < \frac{1}{k}$, then $(\frac{2k}{k+1})x \notin \mathcal{U}(G, r, b)$.*

Proof. Consider a subset S of G' and for $i = 1, 2$ let S_i be the copy of S in G_i. Set $S' = S_1 \cup S_2$ in G. We then have:

$$\begin{aligned}x(\delta_G(S')) &= |S_1| + \sum_{i=1,2} x(\delta_{G_i}(S_i)) \\ &= |S| + u(\delta_{G'}(S)) \qquad (2) \\ &= |S| + u(\delta_{G'}(V' - S)). \qquad (3)\end{aligned}$$

Equation (3) follows from Equation (2) since G' is undirected.

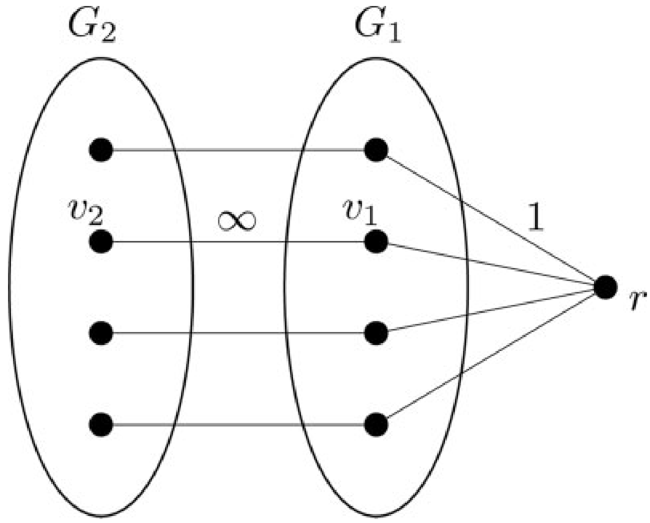

FIG. 1. Reduction from expansion.

Now, suppose S is a light subset of G' with expansion less than $\frac{1}{k}$. Since S was light, there is some demand matrix in $P(G, r, b)$ for which r wishes to send $|S'|$ units of flow to nodes in S'. By (2) $x(\delta_G(S')) < |S| + \frac{1}{k}|S| = (\frac{k+1}{k})|S'|/2$, and so even $\frac{2k}{k+1}x \notin \mathcal{U}(G, r, b)$, as required.

Conversely, suppose that $exp(G') \geq 1$ but $x \notin \mathcal{U}(G, r, b)$. Then there is some demand matrix that cannot be supported. By Fact 2.1, there is an integral such demand, where r is required to send one unit of flow to each node in a set Z of at most $b_r \leq |V'|$ nodes in $G-r$. Suppose S' is a subset of $G-r$, that does not have enough capacity to support this demand. That is, $x(\delta_G(S')) < |Z \cap S'|$. Since this cut has finite capacity, we have that $S' = S_1 \cup S_2$, where for $i = 1, 2$, S_i is the copy of a set S of nodes from G'. Suppose first that S is light in G'. By (2) we have that $x(\delta_G(S')) = |S| + u(\delta_{G'}(S)) < |S' \cap Z| \leq 2|S|$, which implies that $u(\delta_{G'}(S)) < |S|$, contradicting our assumption. So suppose that S is not light, and hence $V' - S$ is light. Thus, by (3) $x(\delta_G(S')) = |S| + u(\delta_{G'}(V' - S)) < |S' \cap Z| \leq |Z| \leq |V'|$. Hence, $u(\delta_{G'}(V' - S)) < |V' - S|$, a contradiction to our assumption that $exp(G') \geq 1$. This completes the proof. ∎

We are now ready to prove the main result.

Theorem 2.3. *For any integer $k \geq 1$, if there is a polynomial time $2\frac{k}{k+1}$-separation algorithm for the single-source hose problem, then there is a polynomial time algorithm with the following property: given an input graph G' it outputs YES if $exp(G') \geq 1$, NO if $exp(G') < 1/k$ and otherwise its output is not defined. In particular, taking $k = 1$ we get that the separation problem for the single-source hose problem is coNP-complete. This holds even if we restrict to instances where $b_i \in \{0, 1\}$ for $i \neq r$.*

Proof. Consider the following algorithm that takes the input graph G'. In polynomial time, it constructs the graph G and x as described in the reduction above. It then queries the $(\frac{2k}{k+1})$-approximate separation algorithm to see if $x \in \mathcal{U}(G, r, b)$. If the separation algorithm claims infeasibility of x, the algorithm outputs NO on G', otherwise it outputs YES. Using Claim 2.2, it is straightforward to check that the algorithm behaves correctly if $exp(G') \geq 1$ or if $exp(G') < 1/k$. Further, the algorithm runs in polynomial time. ∎

Theorem 2.3 establishes that if expansion can be shown to be hard to approximate to within a factor of $1/k$, then the separation problem is hard to approximate to within a factor of $2k/(k+1)$. Recently, under a particular complexity theoretic conjecture called the unique games conjecture (due to Khot [16]), authors of [6] and [18] have independently shown that for any fixed $C \geq 1$, there is no polynomial time C-approximation algorithm for graph expansion. In fact, like in most hardness results, a stronger gap result is shown: it is hard to distinguish between instances with expansion β and those with expansion β/C where β is some polynomial time computable function of the input graph. The unique games conjecture states that a certain constraint satisfaction problem is hard to approximate for some parameter settings. This conjecture has been used to prove hardness of approximation for several problems that have resisted previous attempts. We refer the reader to a recent tutorial of Khot on various aspects of the conjecture [17].

We immediately obtain the following corollary.

Corollary 2.4. *If the unique games conjecture is true, for any fixed $\epsilon > 0$, the $(2 - \epsilon)$-separation problem for single-source hose problem is coNP-complete. This holds even if we restrict to instances where $b_r = \frac{1}{2}\sum_{i \neq r} b_i$.*

We remark that the inapproximability result above relies only on the hardness of expansion and might even hold even if the unique games conjecture is false.

We have established hardness for the case $b_r = \delta \sum_{i \neq r} b_i$, where $\delta = \frac{1}{2}$. We close with a brief investigation of the complexity of the problem when $b_r = \delta \sum_{i \neq r} b_i$ for a rational number $\delta \in (0, 1)$. First, let $\delta \in (0, 1/2]$. Suppose $\delta = s/t$, where s, t are integers. We may assume without loss of generality that the graph G', whose expansion we want to evaluate in the proof of Theorem 2.3, is such that $\frac{|V(G')|}{s}$ is integral. In fact, we may replace each node of G' by a complete graph of size s; if we give a sufficiently large capacity to the edges in the complete graph, the expansion of the resulting graph is the same as that of G'.

So assume that $\frac{|V(G')|}{\delta}$ is integral. We modify the construction of the graph in the proof of Theorem 2.3 to show co-NP hardness for this case. Let $G = (V, E)$ be the graph constructed in the proof and let $n = |V| = 2|V(G')| + 1$. Note that $b_r = \frac{1}{2}(n-1) = \frac{1}{2}\sum_{i \neq r} b_i$. We construct a new graph H by adding $(\frac{1}{2\delta} - 1)(n-1)$ extra nodes to G. We know that $\frac{n-1}{2\delta} = \frac{|V(G')|}{\delta}$ is an integer. The new nodes are connected only to r in the form of a star by 0-cost edges and for each extra node i we set $b_i = 1$. In the new graph, we have that $\sum_{i \neq r} b_i = \frac{1}{2\delta}(n-1)$. We leave b_r at $\frac{1}{2}(n-1)$ and hence in the new instance $b_r = \delta \sum_{i \neq r} b_i$. Clearly, an optimal design will set each new edge to capacity 1 at 0-cost, and the only optimization to be done is on the original edges. Hence the hardness for G carries over to the new instance. This establishes that instances with $b_r = \delta \sum_{i \neq r} b_i$ for $\delta \in (0, 1/2]$ are hard.

Now we consider the case when $\delta \in (1/2, 1)$. We use a scaling argument which when combined with Corollary 2.4 shows hardness for this case assuming that the unique games conjecture is true. Suppose we are given an instance of the single-source hose robust design problem for a graph G with $b_r = \frac{1}{2} \sum_{i \neq r} b_i$. We define a new instance of the problem on the same graph G by altering the b values as follows. We set $b'_r = b_r$ and for $i \neq r$ we set $b'_i = \frac{1}{2\delta} b_i$. Observe that $b'_r = b_r = \delta \sum_{i \neq r} b_i = \frac{1}{2} \sum_{i \neq r} b'_i$. It can be seen that a polynomial time separation algorithm for the new instance directly translates into a polynomial time 2δ-separation algorithm for the original instance. Thus for any $\delta < 1$, a separation algorithm for instances with $b_r = \delta \sum_{i \neq r} b_i$ would contradicting the $(2 - \epsilon)$-hardness result of Corollary 2.4.

Finally, we describe a simple approximation algorithm.

Theorem 2.5. *There is a polynomial time* $\max(\frac{1}{\delta}, 1)$-*approximation algorithm for the single-source hose robust design problem, for the class of instances where* $b_r = \delta \sum_{i \neq r} b_i$.

Proof. Let $\alpha = \min(\delta, 1)$. Recall that $P(G, r, b)$ is the polytope of demand matrices and $\mathcal{U}(G, r, b)$ the region $\mathcal{U}(P(G, r, b))$. Let u^* be an optimal reservation vector. Consider the demand matrix $\tilde{D}$ given by: $\tilde{D}_{ri} = \alpha \cdot b_i$ for all $i \neq r$ and $\tilde{D}_{ij} = 0$ if $i, j \neq r$.

The demand matrix $\tilde{D}$ belongs to $P(G, r, b)$, therefore u^* must support it. Let $\tilde{u}$ be a minimum cost capacity reservation supporting $\tilde{D}$: $\tilde{u}$ may be evaluated by computing a shortest path tree rooted at r. Clearly, $c \cdot \tilde{u} \leq c \cdot u^*$. Observe that $\frac{1}{\alpha}\tilde{u} \in \mathcal{U}(P(G, r, b))$ and therefore $\frac{1}{\alpha}\tilde{u}$ is our $\frac{1}{\alpha}$-approximate solution. ∎

Note that the proof of Theorem 2.3 actually holds even for the restricted version where b_r is (about) one half of $\sum_{i \neq r} b_i$; therefore in the case where $\delta = \frac{1}{2}$, the inapproximability bound of Corollary 2.4 is tight. We mention that we presently do not know whether there is a constant approximation when the value of δ is small.

REFERENCES

[1] A. Altin, E. Amaldi, P. Belotti, and M.C. Pinar, Provisioning virtual private networks under traffic uncertainty, Networks 49 (2007), 100–115.

[2] D. Applegate and E. Cohen, Making intra-domain routing robust to changing and uncertain traffic demands: Understanding fundamental tradeoffs, Proc SIGCOMM, 2003.

[3] Y. Azar, E. Cohen, A. Fiat, H. Kaplan, and H. Räcke, Optimal oblivious routing in polynomial time, J Comput Syst Sci 69 (2004), 383–394. Preliminary version in *Proc ACM STOC*, 2003.

[4] W. Ben-Ameur and H. Kerivin, Routing of uncertain demands, Optimization Eng 3 (2005), 283–313.

[5] W. Ben-Ameur and H. Kerivin, New economical virtual private networks, Commun ACM 46 (2003), 69–73.

[6] S. Chawla, R. Krauthgamer, Ravi Kumar, Y. Rabani, and D. Sivakumar, On the hardness of approximating multicut and sparsest-cut, Proc IEEE Conf Comput Complexity (2005), 144–153.

[7] N.G. Duffield, P. Goyal, A. Greenberg, P. Mishra, K.K. Ramakrishnan, and J.E. van der Merwe, A flexible model for resource management in virtual private networks, Proc SIGCOMM, 1999.

[8] F. Eisenbrand, F. Grandoni, G. Oriolo, and M. Skutella, New approaches for virtual private network design, Proc ICALP (2005), 1151–1162.

[9] T. Erlebach and M. Rueegg, Optimal bandwidth reservation in hose-model VPNs with multi-path routing, Proc INFOCOM, March 2004.

[10] J. Andrew Fingerhut, S. Suri, and J. Turner, Designing least-cost nonblocking broadband networks, J Algorithms 24 (1997), 287–309.

[11] M.R. Garey, D.S. Johnson, and L. Stockmeyer, Some simplified NP-complete graph problems, Theor Comput Sci 1 (1976), 237–267.

[12] M. Grotschel, L. Lovasz, and A. Schrijver, Geometric algorithms and combinatorial optimization, Springer, Berlin, 1988.

[13] A. Gupta, J. Kleinberg, and T. Roughgarden, Simpler and better approximation algorithms for network design, Proc ACM STOC (2003), 365–372.

[14] A. Gupta, J. Kleinberg, A. Kumar, R. Rastogi, and B. Yener, Provisioning a virtual private network: A network design problem for multicommodity flows, Proc ACM STOC (2001), 389–398.

[15] G. Italiano, S. Leonardi, and G. Oriolo, Design of networks in the hose model, Proc ARACNE (2002), 65–76.

[16] S. Khot, On the power of unique 2-prover 1-round games, Proc ACM STOC (2002), 767–775.

[17] S. Khot, On the unique games conjecture, Tutorial at IEEE FOCS, 2005. Slides available at http://www.cc.gatech.edu/~khot/papers/FOCS05.ppt.

[18] S. Khot and N. Vishnoi, The unique games conjecture, integrality gap for cut problems and embeddability of negative type metrics into ℓ_1, Proc IEEE FOCS (2005), 53–62.

[19] C. Hurkens, J. Keijsper, and Leen Stougie, Virtual private network design: A proof of the tree routing conjecture on ring networks, Proc IPCO (2005), 407–421.

[20] M. Kodialam, T. Lakshman, and S. Sengupta, Maximum throughput routing of traffic in the hose model, INFOCOM, to appear.

[21] T. Leighton and S. Rao, Multicommodity max-flow min-cut theorems and their use in designing approximation algorithms, JACM 46 (1999), 787–832. Preliminary version appeared in Proc IEEE FOCS, 1988.

[22] C. Lemaréchal, A. Ouorou, and G. Petrou, Robust network design in telecommunications, Proc INOC (2005), 258–264.

[23] P. Oswald, F.B. Shepherd, P. Winzer, and M. Zirngibl, Robust network design and selective randomized load balancing, The 31st European Conference on Optical Communications (ECOC 2005), SECC, Glasgow, September 2005, pp. 25–29.

[24] H. Räcke, Minimizing congestion in general networks, Proc IEEE FOCS (2002), 43–52.